

High Performance Distributed Discovery System

5 Field of the Invention

The present invention relates to the discovery of the network topology of devices comprising a high speed data network, and more particularly to a high performance distributed discovery system.

10 Background of the Invention

Today's high speed data networks contain an ever-growing number of devices. A network needs to be monitored for the existence, disappearance, reappearance and status of traditional network devices such as routers, hubs and bridges and more recently high speed switching devices such as ATM, Frame Relay, DSL, VoIP and Cable Modems.

In order to enable network monitoring, a process known as discovery is typically performed. Discovery is the process by which network management systems selectively poll a network to discover very large numbers of objects in a very short period of time, without introducing excessive network traffic. It is the function of a discovery system to discover devices on a network and the structure of that network. Discovery is primarily intended to get network management users quickly up to speed, track changes in the network, update network maps, and report on these changes.

Discovery typically further involves discovering the configuration of individual devices, their relationship, as well as discovering interconnection links or implied relationships.

In the past rapid discovery was not an issue, since the level of scalability of performance monitoring did not require the depth of discovery that is now required. Major advances in scalability have recently been achieved in performance monitoring,

and as performance monitoring scales to manage larger and larger networks the scalability of discovery must advance accordingly in order to deal with the inevitable increase in the number of network objects and react quickly to changes in network topology.

5

At present network devices are typically polled over long distances from the network management system. This consumes valuable bandwidth and results in increased processing times and potential data loss. As well, customers often dislike inadvertent access around their firewalls, via the common connection to the network performance monitoring server computer. Therefore, what is needed is a method of object discovery that is proximal to the managed network.

10

For the foregoing reasons, there is a need for an economical method of network topology discovery that provides for high speed polling, high object capacity, scalability, and proximity to managed networks, while preserving security policies that are inherent in the network domain configuration.

15

Summary of the Invention

20

The present invention is directed to a high performance distributed discovery system that satisfies this need. The system, leveraging the functionality of a high speed communications network, comprises distributing records of discovered network devices using a plurality of discovery engine instances located on at least one data collection node computer whereby the resulting distributed record compilation comprises a distributed network topology database. The distributed network topology database is accessed using at least one performance monitor server computer to facilitate network management.

25

At least one discovery engine instance is located on the data collection node computers on a ratio of one engine instance to one central processing unit whereby the total number of engine instances is at least two so as to enable the parallel processing of the distributed network topology database.

30

09843471.042601

In aspects of the invention a vendor specific discovery subroutine is launched upon detection by the system of a non-MIB II standard device so as to query the vendor's private MIB using a vendor specific algorithm.

5

Advances in overall scalability are achieved by dividing the workload of network topology discovery across several computing nodes. The discovery job is distributed across all the data collectors such that the only requirement for each data collector is to be able to reach, typically via TCP/IP and SNMP, the nodes and networks for which it is responsible. This reachability requirement already exists for telemetry, in any case, and has therefore already been provided for.

10

Other aspects and features of the present invention will become apparent to those ordinarily skilled in the art upon review of the following description of specific embodiments of the invention in conjunction with the accompanying figures.

15

Brief Description of the Drawings

These and other features, aspects, and advantages of the present invention will become better understood with regard to the following description, appended claims, and accompanying drawings where:

20

Figure 1 is a schematic overview of the high performance distributed discovery system.

Detailed Description of the Presently Preferred Embodiment

25

As shown in figure 1, the high performance distributed discovery system, leveraging the functionality of a high speed communications network 14, comprises at least one data collection (DC) node computer 12 and at least one performance monitor (PM) server computer 18 in network 14 contact with the DC node computers 12.

30

The DC node computers 12 poll and register managed network 14 objects with the resulting distributed record compilation forming a distributed network topology database 16 that is accessed by the PM server computers 18.

5 A plurality of discovery engine instances 20 are located on the DC node computers 12 on a ratio of one engine instance 20 to one central processing unit so as to provide for the parallel processing of the distributed network topology database 16.

10 The discovery engine 20 is comprised of a base program and a scalable family of vendor-specific discovery subroutines. The base program is designed to query and register any IP device and subsequently obtain detailed device, state and topology information for any IP device that responds to an SNMP query, such as any device that is managed by an SNMP agent. The base program discovers detailed information for any device that supports the standard MIB-II, but not the vendor's private MIB. The discovery
15 of detailed information from a vendor's private MIB is accomplished through what is known as vendor-specific discovery subroutines.

20 These discovery subroutines are lightweight independent applications that are launched whenever the main discovery program detects a particular vendor's hardware. The discovery subroutines contain vendor-specific algorithms designed to query the vendor's private MIB.

25 Launch points for each discovery subroutine are included in the main program. So, if during the normal operation of discovery a valid element value is encountered identifying a specific vendor's hardware, the appropriate discovery subroutine is launched.

30 The DC node computers 12 are responsible for telemetry to the managed elements and management of the topology database 16. The PM server computers 18 provide system control and reporting interface.

The proximal topology of the DC node computers 12 in relation to the managed network 14 provides for inherent scalability and a reduction in required bandwidth. As well, the ability to utilize excess memory and disk storage resources on the DC node computers 12 facilitates the discovery of larger networks. The aggregate resources of many DC node computers 12 is far greater than that available on any one PM server computer 18. Advances in overall scalability are achieved by dividing the workload of network topology discovery across several computing nodes. The discovery job is distributed across all the DC node computers 12 such that the only requirement for each DC node computer 12 is to be able to reach, typically via TCP/IP and SNMP, the nodes and networks for which it is responsible. This reachability requirement already exists for telemetry, in any case, and has therefore already been provided for.

All the discovery and topology database storage is taking place behind the client's firewall requiring only a minimal amount of management traffic to be exerted on the network to generate reports. PM server computers 18 are utilized to access the distributed network topology database 16 for object management.

In embodiments of the invention unique algorithms selectively discover network devices based on "clues" picked up from existing information such as router tables and customer input.

The vendor specific discovery subroutines extend the base discovery application to provide for inter-operability with a multiplicity of ATM and FR vendors' equipment.

All of the processing intensive data collection takes place as close to the customers network and network devices as possible, thereby providing for faster discovery as well as distributed storage and processing. As well, the unwanted side-effect of the PM server computer 18 unwittingly becoming a router is removed, thereby enhancing security.

Devices are reliably re-discovered, thereby enabling the tracking of changes to a network's topology as it evolves in real time or near real time.

5 The ability to limit what is discovered by criteria such as vendor & device type has been added thereby eliminating the need to specify the address of each device when discovering the network.

10 The system will not re-discover existing devices unless explicitly requested to do so, which is significant when discovering a large network that is typically discovered in stages.

The system handles timeouts in a more reliable manner. This is important on wide area networks where timeouts are more common during discovery.

15 Since all the discovery sub-tasks can be performed simultaneously, the overall time to characterize the customer's network is reduced. This enables discovery to deal with larger networks in a faster manner, and eliminates the PM server computer's 18 reachability requirement with respect to managed elements.

20 This invention allows Network Service Providers to automatically discover more of the existing devices in their networks, permitting customers to reconcile what is really out in their network with what their administrative records tell them is out there. It has been shown that such verification can potentially lead to great cost savings in operations, as well as vastly improved discovery times as speed will now be directly correlated with
25 the number of DC node computers 12 deployed.

The system provides for the rapid automatic mapping of a customer's network for the purpose of object management, down to unprecedentedly fine levels of granularity.